PAPER

Monitoring Temporal Properties using Interval Analysis*

Daisuke ISHII†, **Naoki YONEZAKI**†, ***Members***, ***and*** **Alexandre GOLDSZTEJN**††, ***Nonmember***

arXiv:1602.02387v1 [cs.LO] 7 Feb 2016

SUMMARY Verification of temporal logic properties plays a crucial role in proving the desired behaviors of continuous systems. In this paper, we propose an interval method that verifies the properties described by a bounded signal temporal logic. We relax the problem so that if the verification process cannot succeed at the prescribed precision, it outputs an inconclusive result. The problem is solved by an efficient and rigorous monitoring algorithm. This algorithm performs a forward simulation of a continuous-time dynamical system, detects a set of time intervals in which the atomic propositions hold, and validates the property by propagating the time intervals. In each step, the continuous state at a certain time is enclosed by an interval vector that is proven to contain a unique solution. We experimentally demonstrate the utility of the proposed method in formal analysis of nonlinear and complex continuous systems.

key words: *continuous-time dynamical systems, interval analysis, linear temporal logic, falsification method*

1. Introduction

Reasoning with the temporal logic properties in continuous systems is a challenging and important task that combines computer science, numerical analysis, and control theory. Various methods for the verification of continuous and hybrid systems with bounded temporal properties have been developed, e.g., [1]–[4], enabling the falsification of various properties (e.g., safety, stability, and robustness) of large and complex systems. However, the state-of-the-art tools are based on numerical simulations whose numerical errors frequently yield qualitatively wrong results, which become problematic even in statistical evaluations.

Computing rigorously approximated reachable sets is a fundamental process in formal methods for continuous systems. Techniques based on interval analysis (Section 3) have proven practical in the reachability analysis of nonlinear and complex continuous systems [5]–[10]. In these frameworks, the computation is δ-*complete* [11]: assuming that function values may be perturbed within a predefined $\delta \in \mathbb{Q}_{>0}$, many generically undecidable problems become decidable. However, δ-complete verification of generic properties other than reachability is a challenging topic.

The contribution of this paper is to propose an interval method that verifies (bounded portions of) the signal temporal logic (STL) properties (Section 5) of a class of continuous-time dynamical systems (Section 4; extension to hybrid systems is straightforward). Our method reliably computes three values: valid, unsat, and unknown. The method outputs valid or unsat when the soundness is guaranteed by interval analysis; otherwise, when the verification fails after reaching a prescribed precision threshold, it outputs unknown. Our method is based on validated interval analysis, and therefore it is reliable compared to the existing simulation-based monitoring tools, e.g., [12]–[15]. We show that simulation with numerical errors may compute an incorrect signal for a chaotic system. In contrast with the existing tools that monitor a single behavior of a system, our method monitors a set of possible behaviors using an interval-based technique; therefore, the method can check the validity of the system. In this sense, our approach can be viewed as an integration of reachability analysis and simulation-based monitoring methods. The relaxation allowing unknown results enables us to generate an efficient monitor for STL properties that can be regarded as a variant of δ-complete procedures. We demonstrate efficient and reliable monitors for several continuous systems including a chaotic system.

In Section 6, we present an algorithm for monitoring STL properties based on the forward simulation that encloses a signal with a set of *boxes* (i.e., interval vectors). For each atomic proposition involved in a property φ to be verified, the algorithm obtains an inner and outer approximation of the time intervals in which the proposition holds. The interval Newton operator is used for the purposes of accelerating the search of instants where the satisfaction of propositions changes, and certifying the uniqueness of these event within their enclosures, eventually certifying the sequence of consistent/inconsistent time intervals over time for each proposition. Next, it modifies the set of time intervals according to the syntax of the property φ; finally, it checks that φ holds at the initial time. Using our implementation, we show that several benchmarks are verified efficiently, yet non-robust instances with respect to numerical errors are rejected (Section 7). The implementation reliably analyzes a set of signals and

Manuscript received August 11, 2015.
Manuscript revised October 5, 2015.

†The authors are with Tokyo Institute of Technology.
††The author is with CNRS/IRCCyN.

*The preliminary version of this paper was presented at the 8th International Workshop on Numerical Software Verification (NSV'15).

DOI: 10.1587/transfun.E99.A.1

Copyright © 2016 The Institute of Electronics, Information and Communication Engineers

provides a foundation for verification and parameter synthesis of complex systems.

2. Related Work

Many previous studies have applied interval methods to reachability analyses of continuous and hybrid systems [5]–[10]. These methods output an over-approximation of reachable states as a set of boxes. Interval analysis often proves the unique existence of a solution within a resulting interval, and it is also applicable to interval-based reachability analysis [8], [16]. Our method utilizes the proof in the verification of more generic temporal properties.

Reasoning of real-time temporal logic has been a research topic of interest [17], [18]. Numerical method for *falsification* of a temporal property is straightforward [12]. The algorithm simulates a signal of a bounded length and checks the satisfiability of the negation of the property described by a bounded temporal logic. This paper presents an interval extension of this falsification method.

To falsify realistic nonlinear models efficiently, researchers have proposed a tree-search method [1], a Monte-Carlo optimization method [2], and statistical model checking methods [3], [4]. Despite their successes, these methods are compromised by numerical error. To improve the reliability and practicality of the falsification, integration with our interval method will be a promising future direction. An integrated statistical and interval method was also proposed in [19] for reachability analysis.

To facilitate simulation-based verification of temporal properties, the *robustness* concept has been proposed [2], [13], [14]. In these works, the degree of robustness defines the distance between a signal and a region over which a proposition holds. If the absolute value of the degree is small, it is likely to be unreliable because of numerical errors. Our method rigorously ensures robustness by verifying that every intersection between a signal and each boundary in the state space is enclosed with an interval.

There exist several methods for model checking of temporal logic properties [20], [21]. [20] proposed a method specialized in stability properties, which is described as a specific form of temporal logic formula. [21] proposed a method that translates a verification problem into a reachability problem with the k-Liveness scheme, which is incomplete in general settings. Our method can be viewed as a bounded model checking method that validates a bounded temporal property when the property is satisfied by all signals emerging from the interval parameter value.

3. Interval Analysis

This section introduces selected topics and techniques based on interval analysis [22], [23].

A (bounded) *interval* $\boldsymbol{a} = [\underline{a}, \overline{a}]$ is a connected set of real numbers $\{b \in \mathbb{R} \mid \underline{a} \leq b \leq \overline{a}\}$. $\mathbb{I}$ denotes the set of intervals. $\mathbb{I}_{\geq 0}$ denotes the subset $\{[\underline{a}, \overline{a}] \in \mathbb{I} \mid \underline{a} \leq 0\}$. For an interval $\boldsymbol{a}$, $\underline{a}$ and $\overline{a}$ denote the lower and upper bounds, respectively; and $\operatorname{int} \boldsymbol{a}$ denotes the interior $\{b \in \mathbb{R} \mid \underline{a} < b < \overline{a}\}$. $[a]$ denotes a point interval $[a, a]$. The hypermetric between two intervals $\boldsymbol{a}$ and $\boldsymbol{b}$, $d(\boldsymbol{a}, \boldsymbol{b})$ is given by $\max(|\overline{a} - \overline{b}|, |\underline{a} - \underline{b}|)$. For a set $S \subset \mathbb{R}$, $\square S$ denotes the interval $[\inf S, \sup S]$. All these definitions are naturally extended to interval vectors; an n-dimensional *box* (or interval vector) $\boldsymbol{a}$ is a tuple of n intervals $(\boldsymbol{a}_1, \ldots, \boldsymbol{a}_n)$, and $\mathbb{I}^n$ denotes the set of n-dimensional boxes. For $a \in \mathbb{R}^n$ and $\boldsymbol{a} \in \mathbb{I}^n$, we use the notation $a \in \boldsymbol{a}$, which is interpreted as $\forall i \in \{1, \ldots, n\}\ a_i \in \boldsymbol{a}_i$.

In actual implementations, the interval bounds should be machine-representable floating-point numbers, and other real values are rounded in the appropriate directions.

Given a function $f : \mathbb{R}^n \to \mathbb{R}$, $\boldsymbol{f} : \mathbb{I}^n \to \mathbb{I}$ is called an *interval extension* of f if and only if it satisfies the containment condition $\forall \boldsymbol{a} \in \mathbb{I}^n\ \forall a \in \boldsymbol{a}\ (f(a) \in \boldsymbol{f}(\boldsymbol{a}))$. This definition is generalizable to function vectors $\boldsymbol{f} : \mathbb{R}^n \to \mathbb{R}^m$. Given two intervals $\boldsymbol{a}, \boldsymbol{b} \in \mathbb{I}$, we can compute interval extensions of the four operators $\circ \in \{+, -, *, /\}$ as $\square\{\underline{a} \circ \underline{b}, \underline{a} \circ \overline{b}, \overline{a} \circ \underline{b}, \overline{a} \circ \overline{b}\}$ (assuming $0 \notin \boldsymbol{b}$ for division).

For arbitrary intervals $\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{d} \in \mathbb{I}$, the *extended division* $\square\{d \in \boldsymbol{d} \mid \exists a \in \boldsymbol{a}\ \exists b \in \boldsymbol{b}\ a = bd\}$ can be implemented as follows (see Section 4.3 of [23]):

$$\mathsf{ExtDiv}(\boldsymbol{a}, \boldsymbol{b}, \boldsymbol{d}) := \begin{cases} (\boldsymbol{a}/\boldsymbol{b}) \cap \boldsymbol{d} & \text{if } 0 \notin \boldsymbol{b} \\ \square(\boldsymbol{d} \setminus (\underline{a}/\underline{b}, \underline{a}/\overline{b})) & \text{if } \underline{a} > 0 \in \boldsymbol{b} \\ \square(\boldsymbol{d} \setminus (\overline{a}/\overline{b}, \overline{a}/\underline{b})) & \text{if } \overline{a} < 0 \in \boldsymbol{b} \\ \boldsymbol{d} & \text{if } 0 \in \boldsymbol{a}, \boldsymbol{b} \end{cases}$$

In the second and third cases, when $\underline{b} = 0$ (resp. $\overline{b} = 0$), we set $\underline{a}/\underline{b}$ and $\overline{a}/\underline{b}$ as $-\infty$ and ∞ (resp. $\underline{a}/\overline{b}$ and $\overline{a}/\overline{b}$ as ∞ and $-\infty$).

Given a differentiable function $f(a) : \mathbb{R} \to \mathbb{R}$ and a domain interval $\boldsymbol{a}$, a root $\tilde{a} \in \boldsymbol{a}$ of f such that $f(\tilde{a}) = 0$ is included in the result of the *interval Newton operator*

$$\hat{a} + \mathsf{ExtDiv}(-\boldsymbol{f}(\hat{a}), \boldsymbol{f}'(\boldsymbol{a}), \boldsymbol{a} - \hat{a}) \approx \left(\hat{a} - \frac{\boldsymbol{f}(\hat{a})}{\boldsymbol{f}'(\boldsymbol{a})}\right) \cap \boldsymbol{a},$$

where $\hat{a} \in \boldsymbol{a}$, and $\boldsymbol{f}$ and $\boldsymbol{f}'$ are interval extensions of f and the derivative of f, respectively. The first expression is always valid while the second expression is valid only when $\boldsymbol{f}'(\boldsymbol{a})$ does not contain 0. Iterative applications of the operator will converge. Let $\boldsymbol{a}'$ be the result of applying the operator to $\boldsymbol{a}$. If $\boldsymbol{a}' \subseteq \operatorname{int} \boldsymbol{a}$, a unique root exists in $\boldsymbol{a}'$.

4. Continuous-Time Dynamical Systems

We consider dynamical systems whose behaviors are

described by ordinary differential equations (ODEs).

Definition 1 *A* continuous-time dynamical system *is a tuple* $\mathcal{S} := \big((u,x), U\times X, X_{\text{init}}, F\big)$ *consisting of the following components:*

- *A vector of real-valued* parameters $u = (u_1,\ldots,u_m)$.
- *A vector of real-valued* variables $x = (x_1,\ldots,x_n)$.
- *A* domain $U\times X \subseteq \mathbb{R}^{m+n}$ *for the valuation of the parameters and variables.*
- *An* initial domain $X_{\text{init}} \subseteq X$.
- *A* vector fields $F : U\times X \to \mathbb{R}^n$ *(assuming Lipschitz continuity).*

In this work, we specify domains U and X as boxes. The behaviors of a system $\mathcal{S}$ are formalized as *signals*.

Definition 2 *Given a time interval* $\boldsymbol{t} = [0,\overline{t}] \in \mathbb{I}$ *and a parameter value* $\tilde{u} \in U$, *a* signal *of a continuous-time dynamical system* $\mathcal{S}$ *is a function* $\tilde{x} : \boldsymbol{t} \to X$ *such that*

$$\tilde{x}(0) \in X_{\text{init}} \wedge \forall \tilde{t}\in\boldsymbol{t}\ \tfrac{d}{dt}\tilde{x}(\tilde{t}) = F(\tilde{u},\tilde{x}(\tilde{t})).$$

$SS_{\overline{t}}(\mathcal{S})$ denotes the set of signals of $\mathcal{S}$ of length $\overline{t}$.

Example 1 *An anticlockwise rotation of a 2D particle can be modeled as a continuous-time dynamical system:*

$$\begin{aligned} u &:= (u_1), \quad U := ([-0.1, 0.1]),\\ x &:= (x_1,x_2), \quad X := [-10,10]^2,\\ X_{\text{init}} &:= \{(1,0)\},\\ F(u,x) &:= \begin{pmatrix} u_1 & -1\\ 1 & u_1\end{pmatrix}\begin{pmatrix} x_1\\ x_2\end{pmatrix}. \end{aligned}$$

A signal of this example is illustrated in Figure 1. The signal moves on the circle of radius 1 when $u_1 = 0$*; the system is stable when* $u_1 \le 0$ *and is unstable when* $u_1 > 0$.

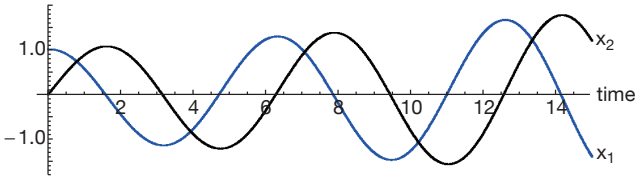

Fig. 1 A signal of the rotation system

Example 2 *A well-known chaotic dynamical system is the Lorenz equation:*

$$\begin{aligned} u &:= (u_1,u_2,u_3),\ U := (10,28,2.5) + [-1,1]^3,\\ x &:= (x_1,x_2,x_3), \quad X := [-50,50]^3,\\ X_{\text{init}} &:= \{(15,15,36)\},\\ F(u,x) &:= \begin{pmatrix} u_1(x_2-x_1)\\ x_1(u_2-x_3)-x_2\\ x_1x_2 - u_3x_3\end{pmatrix}. \end{aligned}$$

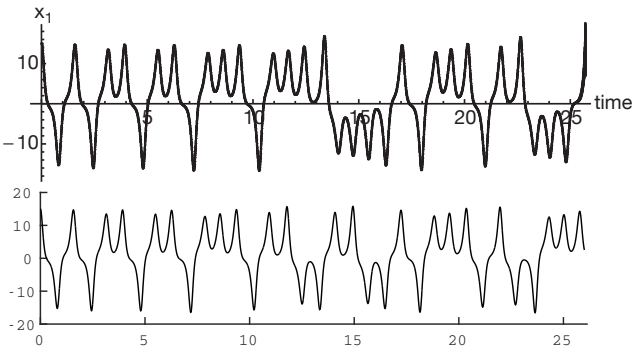

Fig. 2 Signals of the Lorenz system simulated by validated (upper) and non-validated (lower) numerical methods

A signal of this system is illustrated in the upper part of Figure 2.

4.1 ODE Integration using Interval Analysis

Using tools based on interval Taylor methods, such as CAPD[†] and VNODE [24], we can obtain an interval extension $\mathsf{X} : \mathbb{I}_{\ge 0} \to \mathbb{I}^n$ of signals in $SS_{\overline{t}}(\mathcal{S})$. Given $\boldsymbol{t} \in \mathbb{I}_{\ge 0}$, these tools perform stepwise integration of the flow function F from the initial time 0 to time $\overline{t}$, and output the value $\mathsf{X}(\boldsymbol{t})$. At each step, interval Taylor methods verify the *unique existence* of a solution in a box enclosure using the Picard-Lindelöf operator and Banach's fixpoint theorem. Accordingly, when an interval enclosure $\mathsf{X}(\boldsymbol{t})$ is computed by an interval Taylor method, the following property holds:

$$\begin{aligned} &\forall u\in U\ \forall x_{\text{init}}\in X_{\text{init}}\ \exists!\tilde{x}\in SS_{\overline{t}}(\mathcal{S})\\ &\tilde{x}(0) = x_{\text{init}}\ \wedge\ \forall\tilde{t}\in\boldsymbol{t}\ \tfrac{d}{dt}\tilde{x}(\tilde{t}) = F(u,\tilde{x}(\tilde{t})), \end{aligned}$$

where $\exists!$ is interpreted as "uniquely exists."

In principle, if F is Lipschitz continuous and we can assume arbitrary precision, we obtain an arbitrarily narrow interval enclosure $\mathsf{X}([t])$ for $t \in \mathbb{R}_{\ge 0}$. However, because interval Taylor methods are implemented using machine-representable real numbers, they may fail to compute an enclosure when verifying the unique existence property, even at the smallest step size.

Example 3 *Signals of the Lorenz system in Example 2 (when* $u := (10,28,2.5)$*), computed with an interval method (CAPD) and a non-validated numerical method, are illustrated in Figure 2. Non-validated numerical methods may compute a wrong signal for a chaotic system as shown in this figure. On the other hand, validated simulation of this system over a long period is difficult with double-precition floating-point numbers; the width of the interval enclosure computed by CAPD blows up after* 25 *time units and the simulation fails.*

[†] `http://capd.ii.uj.edu.pl/`

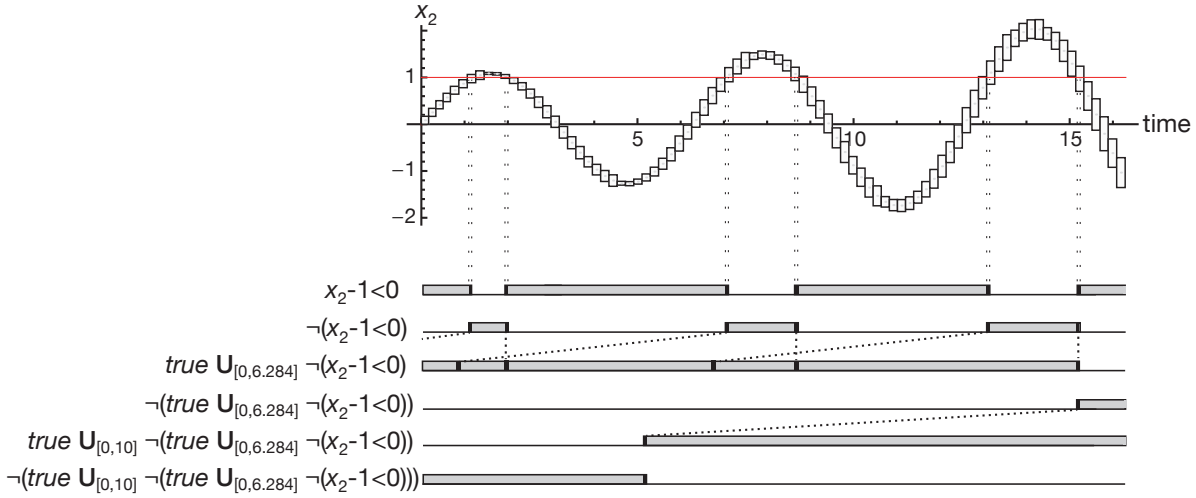

Fig. 3 Monitoring process on the rotation system

5. Signal Temporal Logic

We consider a fragment [12] of the real-time metric temporal logic [17] whose temporal modalities are bounded by an interval $\boldsymbol{t} = [\underline{t}, \overline{t}]$, where the bounds $\underline{t}, \overline{t}$ are in $\mathbb{Q}_{\geq 0}$. Following [12], we refer to this logic as the *signal temporal logic* (STL).

Definition 3 *We consider constraints in the real domain as atomic propositions. The syntax of the STL formulae is defined by the grammar*

$$\varphi ::= \mathsf{true} \mid p \mid \varphi \vee \varphi \mid \neg\varphi \mid \varphi \,\mathsf{U}_{\boldsymbol{t}}\, \varphi$$
$$p ::= f(x) < 0$$

where p belongs to a set of atomic propositions *AP_φ, $\mathsf{U}_{\boldsymbol{t}}$ is the "until" operator bounded by a non-empty positive time interval $\boldsymbol{t} \in \mathbb{I}_{\geq 0}$, $x = (x_1, \ldots, x_n)$ is a vector of variables, and $f : \mathbb{R}^n \to \mathbb{R}$. We use the standard abbreviations, e.g., $\varphi_1 \wedge \varphi_2 := \neg(\neg\varphi_1 \vee \neg\varphi_2)$, $\mathsf{F}_{\boldsymbol{t}}\varphi := \mathsf{true}\,\mathsf{U}_{\boldsymbol{t}}\,\varphi$ ("eventually"), and $\mathsf{G}_{\boldsymbol{t}}\varphi := \neg\mathsf{F}_{\boldsymbol{t}}\neg\varphi$ ("always").*

5.1 Semantics

The necessary length $\|\varphi\|$ of the signals for checking an STL formula φ is inductively defined by the structure of the formula:

$$\|p\| := 0, \qquad \|\varphi_1 \vee \varphi_2\| := \max(\|\varphi_1\|, \|\varphi_2\|),$$
$$\|\neg\varphi\| := \|\varphi\|, \quad \|\varphi_1 \,\mathsf{U}_{\boldsymbol{t}}\, \varphi_2\| := \max(\|\varphi_1\|, \|\varphi_2\|) + \overline{t}.$$

The map $\mathcal{O} : AP_\varphi \to 2^X$ associates each proposition $p \in AP_\varphi$ to a set $\mathcal{O}(p) = \{x \in X \mid p(x)\}$.

When we check the satisfiability of φ at time t, we should have a signal of length $t_{\max} := \|\varphi\| + t$ (this value of $t_{\max}$ is used in evaluating all subformulae of φ). Let $\tilde{x} \in SS_{t_{\max}}(\mathcal{S})$ and φ be an STL property. Then, we have a satisfaction relation defined as follows:

$$\begin{aligned}
&\tilde{x}, t \models \mathsf{true} \\
&\tilde{x}, t \models p && \text{iff } \ \tilde{x}(t) \in \mathcal{O}(p) \\
&\tilde{x}, t \models \varphi_1 \vee \varphi_2 && \text{iff } \ \tilde{x}, t \models \varphi_1 \ \vee \ \tilde{x}, t \models \varphi_2 \\
&\tilde{x}, t \models \neg\varphi && \text{iff } \ \tilde{x}, t \not\models \varphi \\
&\tilde{x}, t \models \varphi_1 \,\mathsf{U}_{\boldsymbol{t}}\, \varphi_2 \\
&\quad \text{iff } \exists t' \in (t + \boldsymbol{t}) \ \tilde{x}, t' \models \varphi_2 \ \wedge \ (\forall t'' \in [t, t'] \ \tilde{x}, t'' \models \varphi_1)
\end{aligned}$$

At a given time t, $\varphi_1 \,\mathsf{U}_{\boldsymbol{t}}\, \varphi_2$ intuitively means that φ_2 holds within the time interval $t + \boldsymbol{t}$ and that φ_1 always hold until then. We also have the following validation relation:

$$\mathcal{S} \models \varphi \ \text{ iff } \ \forall \tilde{x} \in SS_{\|\varphi\|}(\mathcal{S}) \ \tilde{x}, 0 \models \varphi$$

5.2 Method for Monitoring STL Formulae

Our interval method is based on the monitoring method proposed in [12], which decides whether a signal satisfies an STL property based on the numerical simulation of signals of bounded lengths. This section explains this basic method. First, we introduce the notion of consistent time intervals in the STL evaluation.

Definition 4 *Let $\tilde{x}$ be a signal of length $t_{\max}$ and φ be an STL formula. We say that a left-closed and right-open interval $[\underline{t}, \overline{t}) \subseteq \mathbb{R}_{\geq 0}$ is* consistent *with φ iff $\forall t \in (\underline{t}, \overline{t}) \ \tilde{x}, t \models \varphi$.*[†]

[†]The original definition [12] involves left-closed right-open time intervals $[\underline{t}, \overline{t})$ so that they do not overlap and they can cover $[0, t_{\max}]$. However, $\tilde{x}(t) > 1 \ \equiv \ 1 - \tilde{x}(t) < 0$, with $\tilde{x}(t) := t$, is not true in the left-closed right-open interval $[1, 2)$. In this paper, we only enforce the predicate to be true in the interior of time intervals $(\underline{t}, \overline{t})$ to regard $[1, 2)$ consistent. This has no impact on the soundness nor efficiency of the proposed method, since such bounds will be approximated by enclosing intervals in Definition 5.

Next, whether a signal satisfies property φ is checked as follows:

1. For each atomic proposition p in AP_φ, monitor the signal of length $\|\varphi\|$ and identify a non-overlapping set of consistent time intervals $T_p = \{\boldsymbol{t}_1, \ldots, \boldsymbol{t}_{n_p}\}$.
2. Following the parse tree of φ in a bottom-up fashion, obtain a set of consistent time intervals of φ. For each construct of STL, obtain the set that is consistent with the sub-formula as follows:

$$T_{\neg\varphi} := \mathbb{R}_{\geq 0} \setminus T_\varphi \tag{1}$$
$$T_{\varphi_1 \vee \varphi_2} := T_{\varphi_1} \cup T_{\varphi_2} \tag{2}$$
$$T_{\varphi_1 \mathsf{U}_{\boldsymbol{t}} \varphi_2} := \{\mathsf{Shift}_{\boldsymbol{t}}(\boldsymbol{t}_1 \cap \boldsymbol{t}_2) \cap \boldsymbol{t}_1 \mid \boldsymbol{t}_1 \in T_{\varphi_1}, \boldsymbol{t}_2 \in T_{\varphi_2}\} \tag{3}$$

   where $\mathsf{Shift}_{\boldsymbol{t}}(\boldsymbol{s}) := [\underline{s} - \overline{t}, \overline{s} - \underline{t}) \cap \mathbb{R}_{\geq 0}$.
3. Check whether $\min T_\varphi$ contains time 0. If yes, φ is satisfied; otherwise, it is not satisfied.

Example 4 *We consider the property*

$$\mathsf{G}_{[0,10]}\mathsf{F}_{[0,6.284]}\,\neg(x_2 - 1 < 0) \;\equiv\; \neg(\mathsf{true}\,\mathsf{U}_{[0,10]}\,\neg(\mathsf{true}\,\mathsf{U}_{[0,6.284]}\,\neg(x_2 - 1 < 0)))$$

for the model in Example 1, which describes that, within the initial 10 *time units, the signal* x_2 *increases beyond* 1 *within every* 6.284 *time units. Verification with the monitoring method (extended to an interval method) is illustrated in Figure 3, when the parameter is set as* $u_1 := 0.1 + [-10^{-3}, 10^{-3}]$.

6. Interval-Based Monitoring Method

In this section, we propose a reliable method for monitoring STL properties on continuous-time dynamical systems. This method is an interval extension of the monitoring method described in Section 5.2.

Given a system $\mathcal{S}$ and an STL property φ, the proposed MonitorSTL algorithm (Algorithm 1) outputs the following results: valid (implying that $\mathcal{S} \models \varphi$), unsat (implying that $\mathcal{S} \models \neg\varphi$), or unknown (meaning that the computation is inconclusive). The algorithm implements the method described in Section 5.2. The sub-procedures MonitorAP (Section 6.2) for monitoring atomic propositions, and Propagate and ConsistentAtInitTime (Section 6.3) for evaluating an STL formula are rendered rigorous and sound by interval analysis; namely, the precision of every numerical computation is guaranteed, and the correctness of the monitoring method is assured by verifying the unique existence of a solution within its resulting interval. Any errors introduced by the sub-procedures are captured by the catch clause at Line 5.

Despite its efficient computational cost (Section 6.4), the proposed method has some limitations.

Algorithm 1 MonitorSTL algorithm

Input: $\mathcal{S}$, φ
Output: valid, unsat, or unknown

```
1: try
2:     T := MonitorAP(S, φ)          {Step 1; Section 6.2}
3:     T_φ := Propagate(T, φ))       {Step 2; Section 6.3}
4:     return ConsistentAtInitTime(T_φ)   {Step 3}
5: catch error return unknown end try
```

First, the method is incomplete; it allows inconclusive computations and outputs unknown when the interval computation is too imprecise to separate several solutions within an interval. In practice, the unknown output is valuable, because a numerically non-robust signal is rejected as an error in the verification process. Second, although the algorithm validates system properties in principle, its success is guaranteed only over sufficiently small domains U and X_0, particularly when evaluating nonlinear systems. Third, the method is a bounded model-checking method, in the sense that the domain $U \times X$ and the lengths of the signals are both bounded.

Our method is targeted at (i) a more generic framework, in which the possible initial and parameter values can be exhaustively enumerated, and (ii) statistical methods that treat the parameters as random variables and evaluate probabilistic STL properties.

6.1 Approximation of Consistent Time Intervals

In this section, we introduce an interval approximation for the consistent time intervals (Definition 4). The basic idea is to enclose each bound of the consistent time intervals within a closed interval.

Definition 5 *Given a consistent time interval* $\boldsymbol{t} = [\underline{t}, \overline{t}) \subseteq [0, \|\varphi\|)$ *that is consistent with an STL property* φ*, we define an* (interval) approximation *as a pair* $(\boldsymbol{s}, \boldsymbol{s}')$ *such that* $\boldsymbol{s}, \boldsymbol{s}' \in \mathbb{I}$, $\underline{t} \in \boldsymbol{s}$, *and* $\overline{t} \in \boldsymbol{s}'$.

Given an approximation $(\boldsymbol{s}, \boldsymbol{s}')$ and a continuous signal $\tilde{x}$, we have $\forall t \in [\overline{s}, \underline{s}')\ \tilde{x}, t \models \varphi$.

We now approximate a set of consistent time intervals $\{\boldsymbol{t}_1, \ldots, \boldsymbol{t}_{n_\varphi}\}$ as a set (or sequence) of approximations. Instead of the set of pairs $\{(\boldsymbol{s}_1, \boldsymbol{s}'_1), \ldots, (\boldsymbol{s}_{n_\varphi}, \boldsymbol{s}'_{n_\varphi})\}$, we represent a set of approximations with the set $\{(\boldsymbol{s}_1, \mathsf{true}), (\boldsymbol{s}'_1, \mathsf{false}), \ldots, (\boldsymbol{s}_{n_\varphi}, \mathsf{true}), (\boldsymbol{s}'_{n_\varphi}, \mathsf{false})\}$, where the tags true and false represent whether an element corresponds to a lower or an upper bound. A set of approximations is interpreted as both *outer* and *inner approximations*; that is, each consistent time interval $\boldsymbol{t}_i$ is enclosed by the outer approximation $[\underline{s}_i, \overline{s}'_i]$, and the inner approximation $(\overline{s}_i, \underline{s}'_i)$ is contained in $\boldsymbol{t}_i$.

Definition 6 *Consider a set* $\boldsymbol{T} = \{(\boldsymbol{s}_1, b_1), \ldots,$

$(\boldsymbol{s}_{\#\boldsymbol{T}}, b_{\#\boldsymbol{T}})\}$ *where* $\boldsymbol{s}_i \in \mathbb{I}$, $b_i \in \{\mathsf{true}, \mathsf{false}\}$, *and* $\#\boldsymbol{T} \in \mathbb{N}$. *The second element of each pair is a* polarity value *that represents whether the pair is an enclosure of a lower or upper bound of a consistent time interval. We say that* $\boldsymbol{T}$ *is* canonical *iff*

- *the elements can be sorted, i.e.,* $\forall i \in \{1, \ldots, \#\boldsymbol{T}-1\}\ \overline{s}_i < \underline{s}_{i+1}$,
- $\forall i \in \{1, \ldots, \#\boldsymbol{T}\}\ 0 \leq \overline{s}_i$,
- $\forall i \in \{1, \ldots, \#\boldsymbol{T}-1\}\ b_i \neq b_{i+1}$, *and*
- $b_1 = \mathsf{true}$.

We say that $\boldsymbol{T}$ *is an* (interval) approximation *of* $T_\varphi = \{\boldsymbol{t}_1, \ldots, \boldsymbol{t}_{n_\varphi}\}$ *iff*

- $\boldsymbol{T}$ *is canonical,*
- $\forall \boldsymbol{t} \in T_\varphi\ \exists (\boldsymbol{s}, \mathsf{true}) \in \boldsymbol{T}\ \underline{t} \in \boldsymbol{s}$,
- $\forall \boldsymbol{t} \in T_\varphi\ \overline{t} < t_{\max} \Rightarrow \exists (\boldsymbol{s}, \mathsf{false}) \in \boldsymbol{T}\ \overline{t} \in \boldsymbol{s}$,
- $\forall (\boldsymbol{s}, \mathsf{true}) \in \boldsymbol{T}\ \exists! \boldsymbol{t} \in T_\varphi\ \underline{t} \in \boldsymbol{s}$, *and*
- $\forall (\boldsymbol{s}, \mathsf{false}) \in \boldsymbol{T}\ \exists! \boldsymbol{t} \in T_\varphi\ \overline{t} \in \boldsymbol{s}$.

Given a set of consistent time intervals, its canonical approximation is a disjoint sequence of lower and upper bound enclosures; the sequence starts with a lower bound enclosure and ends with either a lower or an upper bound enclosure. For $\boldsymbol{t} \in T_\varphi$ such that $\overline{t} > \|\varphi\|$, $\boldsymbol{T}_\varphi$ may contain only its lower-bound enclosure. $\boldsymbol{T}_{\mathsf{true}} := \{([0], \mathsf{true})\}$ and $\boldsymbol{T}_{\mathsf{false}} := \emptyset$ are the approximations of $T_{\mathsf{true}} = \mathbb{R}_{\geq 0}$ and $T_{\mathsf{false}} = \emptyset$, respectively.

Example 5 *Let* $\mathcal{S}$ *be* $(x, [0, 10], 0, F(x) = 1)$ *such that the variable* x *represents the signal* $\tilde{x}(t) = t$. *Consider a property* $\varphi := \mathsf{F}_{[0,2\pi]}(\cos x < 0 \wedge \sin x < 0)$. $\|\varphi\|$ *is* 2π, *and the set of consistent time intervals within* $[0, 2\pi]$ *is* $T_{\cos x<0} := \{[\frac{\pi}{2}, \frac{3}{2}\pi)\}$, $T_{\sin x<0} := \{[\pi, 2\pi)\}$, *and* $T_\varphi := \{[0, \frac{3}{2}\pi)\}$, *respectively. Then, their approximations are*

$$\boldsymbol{T}_{\cos x<0} := \{([1.57, 1.58], \mathsf{true}), ([4.71, 4.72], \mathsf{false})\},$$
$$\boldsymbol{T}_{\sin x<0} := \{([3.14, 3.15], \mathsf{true}), ([6.28, 6.29], \mathsf{false})\},$$
$$\boldsymbol{T}_\varphi := \{([0], \mathsf{true}), ([4.71, 4.72], \mathsf{false})\}.$$

6.2 Monitoring Atomic Propositions

This section describes the MonitorAP procedure (Algorithm 2) that, given a system $\mathcal{S}$ and an STL property φ, computes a set $\mathcal{T}$ containing an approximated set $\boldsymbol{T}_p$ of consistent time intervals for each $p \in AP_\varphi$.

The outer loop enumerates each atomic proposition p of the form $f(x) < 0$. Lines 3–4 compute the initial polarity by evaluating the proposition at time 0; the set $\boldsymbol{T}_p$ is initialized accordingly. Note that X represents a solving process for the signals in $SS_{\bar{t}}(\mathcal{S})$ (see Section 4.1), which can be regarded as a function $\mathbb{I}_{\geq 0} \to \mathbb{I}^n$. The inner loop searches for bounds at which f changes sign. Line 7 invokes the SearchZero procedure (Algorithm 3), which searches for the *earliest* bound at which p switches consistency within the time interval $\boldsymbol{t}$, and outputs a sharp enclosure of the bound (or $\emptyset$ if there is no solution). This result is stored in the set $\boldsymbol{T}_p$, and $\boldsymbol{T}_p$ is stored in $\mathcal{T}$.

Algorithm 2 MonitorAP algorithm

Input: $\mathcal{S}$, φ
Output: $\mathcal{T}$

1: $\mathcal{T} = \emptyset$
2: **for** $p = f(x) < 0 \in AP_\varphi$ **do**
3: $\quad b := p(\mathsf{X}(0))$
4: $\quad$ **if** b **then** $\boldsymbol{T}_p := \{([0], b)\}$ **else** $\boldsymbol{T}_p := \emptyset$
5: $\quad \boldsymbol{t} := [0, \|\varphi\|]; \quad b := \neg b$
6: $\quad$ **loop**
7: $\quad\quad \boldsymbol{t} := \mathsf{SearchZero}(\mathsf{X}, F, f, \boldsymbol{t})$
8: $\quad\quad$ **if** $\boldsymbol{t} = \emptyset$ **then break end if**
9: $\quad\quad \boldsymbol{T}_p := \boldsymbol{T}_p \cup \{(\boldsymbol{t}, b)\}; \quad \boldsymbol{t} := [\overline{t}, \|\varphi\|]; \quad b := \neg b$
10: $\quad$ **end loop**
11: $\quad \mathcal{T} := \mathcal{T} \cup \{\boldsymbol{T}_p\}$
12: **end for**
13: **return** $\mathcal{T}$

Example 6 *For* φ *in Example 5, MonitorAP computes* $\mathcal{T}$ *as* $\{\boldsymbol{T}_{\cos x_1<0}, \boldsymbol{T}_{\sin x_1<0}\}$.

The evaluation of atomic propositions $f(x) < 0$ switches between true and false at the root of $f : X \to \mathbb{R}$. The intersection between a signal $\tilde{x}(t)$ and a boundary condition $f(x) = 0$ is searched by Algorithm 3. As inputs, this algorithm accepts an interval extension of the signal $\mathsf{X} : \mathbb{I}_{\geq 0} \to \mathbb{I}^n$, a vector field $F : X \to X$, the function f, and a time interval $\boldsymbol{t}_{\text{init}} \in \mathbb{I}_{\geq 0}$ to be searched. The algorithm computes the time interval $\boldsymbol{t} \subseteq \boldsymbol{t}_{\text{init}}$ that encloses the earliest root, i.e.,

$$\boldsymbol{t} = \square\{\ \min\{t \in \boldsymbol{t}_{\text{init}} \mid f(\tilde{x}(t)) = 0\} \mid \forall \tilde{x} \in SS_{\bar{t}_{\text{init}}}(\mathcal{S})\ \}. \quad (4)$$

SearchZero verifies that $\boldsymbol{t}$ encloses a unique bound, i.e.,

$$\forall \tilde{x} \in SS_{\bar{t}_{\text{init}}}(\mathcal{S})\ \exists! t \in \boldsymbol{t}\ f(\tilde{x}(t)) = 0. \quad (5)$$

Alternatively, if no bound exists in $\boldsymbol{t}_{\text{init}}$, SearchZero verifies the following:

$$\forall \tilde{x} \in SS_{\bar{t}_{\text{init}}}(\mathcal{S})\ \forall t \in \boldsymbol{t}_{\text{init}}\ f(\tilde{x}(t)) \neq 0. \quad (6)$$

Theorem 1 (Soundness) *If SearchZero returns an interval* $\boldsymbol{t} \neq \emptyset$, *properties* (4) *and* (5) *hold. If it returns* $\emptyset$, *property* (6) *holds.*

To justify the soundness os SearchZero, we describe some details of Algorithm 3. Lines 2–6 repeatedly filter the time interval $\boldsymbol{t}$ using the interval Newton operator. Line 4 (and Line 10) invokes the Dt procedure, which is given a function f and computes an interval enclosure of the derivative $\frac{d}{dt} f(\tilde{x}(t))$ over $\boldsymbol{t}$ using the chain rule

$$\tfrac{d}{dt} f(\tilde{x}(\boldsymbol{t})) = \tfrac{d}{dx} f(\tilde{x}(\boldsymbol{t})) \cdot \tfrac{d}{dt} \tilde{x}(\boldsymbol{t}) \subseteq \boldsymbol{f}'(\mathsf{X}(\boldsymbol{t})) \cdot \boldsymbol{F}(\mathsf{X}(\boldsymbol{t})).$$

Algorithm 3 SearchZero algorithm

Input: $\mathsf{X} : \mathbb{I}_{\geq 0} \to \mathbb{I}^n$, $F : X \to X$, $f : X \to \mathbb{R}$, $\boldsymbol{t}_{\text{init}} \in \mathbb{I}_{\geq 0}$
Output: $\boldsymbol{t} \in \mathbb{I}$
Parameter: $\epsilon \in \mathbb{Q}_{>0}$, $\theta \in (0,1) \subset \mathbb{Q}$

1: $\boldsymbol{t} := \boldsymbol{t}_{\text{init}}$
2: **repeat** {Lower bound reduction}
3: $\quad \boldsymbol{t}_{\text{bak}} := \boldsymbol{t}$
4: $\quad \boldsymbol{d} := \mathsf{Dt}(f, \mathsf{X}, F, \boldsymbol{t})$
5: $\quad \boldsymbol{t} := \underline{t} + \mathsf{ExtDiv}(-\boldsymbol{f}(\mathsf{X}(\underline{t})),\ \boldsymbol{d},\ \boldsymbol{t} - \underline{t})$
6: **until** $d(\boldsymbol{t}_{\text{bak}}, \boldsymbol{t}) \leq \epsilon$
7: **if** $\boldsymbol{t} = \emptyset$ **then return** $\emptyset$ **end if**
8: $\boldsymbol{t} := \underline{t}$; $\quad \Delta := \infty$
9: **loop** {Unique solution existence verification}
10: $\quad \boldsymbol{d} := \mathsf{Dt}(f, \mathsf{X}, F, \boldsymbol{t})$
11: $\quad$ **if** $\boldsymbol{d} \ni 0$ **then error end if**
12: $\quad \boldsymbol{t}' := \underline{t} - \boldsymbol{f}(\mathsf{X}(\underline{t}))/\boldsymbol{d}$
13: $\quad$ **if** $\boldsymbol{t}' \subseteq \operatorname{int} \boldsymbol{t}$ **then** $\boldsymbol{t} := \boldsymbol{t}'$; **break end if**
14: $\quad \Delta_{\text{bak}} := \Delta$; $\quad \Delta := d(\boldsymbol{t}, \boldsymbol{t}')$
15: $\quad \boldsymbol{t} := \boldsymbol{t}_{\text{init}} \cap \mathsf{Inflate}(\boldsymbol{t}', 1+\theta)$
16: $\quad$ **if** $\Delta \geq (1-\theta)\,\Delta_{\text{bak}}$ **then error end if**
17: **end loop**
18: **return** $\boldsymbol{t}$

Next, at Line 5, the interval Newton operator is applied. To handle the numerator interval $\boldsymbol{d}$ containing zero, we implement the interval Newton by the extended division described in Section 3. Because we expand the interval Newton on the lower bound $\underline{t}$ and the extended division encloses the values in the $\boldsymbol{t} - \underline{t}$ domain, the resulting $\boldsymbol{t}$ is filtered its inconsistent portion without losing the solutions or being expanded. If the interval Newton returns $\emptyset$, SearchZero also returns $\emptyset$ to signal the unsatisfiability (Line 7).

Because $\boldsymbol{t}$ may contain several solutions, Line 8 of the algorithm resets $\boldsymbol{t}$ to the lower bound as a starting value for computing the enclosure of the earliest solution. Then, SearchZero checks that the time interval contains a unique solution. To this end, it applies the interval Newton with the inclusion test to prove the unique existence of a solution within the contracted interval $\boldsymbol{t}'$ (Lines 9–17). The interval Newton verification is repeated with an inflation process of the time interval (see [25] for a detailed implementation). If Line 18 is reached with no error, the time interval $\boldsymbol{t}$ is a sharp enclosure of the first zero of $f(\tilde{x}(t)) = 0$.

When SearchZero is implemented with machine-representable real numbers or when there is a tangency between the signal and the boundary condition, an error may result. Line 11 of SearchZero outputs an error if the derivative on an (inflated) time interval contains zero. At Line 16, we limit the number of iterations by specifying a threshold $1-\theta$ for the inflation ratio between two consecutive contraction amounts as in [25].

Algorithm 4 Propagate algorithm

Input: φ, $\mathcal{T} = \{\boldsymbol{T}_p\}_{p \in AP_\varphi}$
Output: $\boldsymbol{T}_\varphi$

1: **switch** φ
2: **case** p :
3: $\quad$ **return** $\boldsymbol{T}_p$
4: **case** $\neg\varphi'$:
5: $\quad$ **return** $\mathsf{Invert}(\mathsf{Propagate}(\varphi', \mathcal{T}))$
6: **case** $\varphi_1 \vee \varphi_2$:
7: $\quad \boldsymbol{T}_1 := \mathsf{Propagate}(\varphi_1, \mathcal{T})$
8: $\quad \boldsymbol{T}_2 := \mathsf{Propagate}(\varphi_2, \mathcal{T})$
9: $\quad$ **return** $\mathsf{Join}(\boldsymbol{T}_1, \boldsymbol{T}_2)$
10: **case** $\varphi_1 \mathsf{U}_{\boldsymbol{t}} \varphi_2$:
11: $\quad \boldsymbol{T}_1 := \mathsf{Propagate}(\varphi_1, \mathcal{T})$
12: $\quad \boldsymbol{T}_2 := \mathsf{Propagate}(\varphi_2, \mathcal{T})$
13: $\quad$ **return** $\mathsf{ShiftAll}_{\boldsymbol{t}}(\boldsymbol{T}_1, \boldsymbol{T}_2)$
14: **end switch**

6.3 Evaluation of STL Properties

We now describe the procedures for evaluating STL formulae at Lines 3 and 4 of Algorithm 1. Propagation of a set of monitored time intervals that are consistent with the atomic propositions is implemented as a rigorous and sound but incomplete procedure.

To evaluate the approximated sets, we extend the evaluation procedure on sets of consistent time intervals described in Section 5.2. Algorithm 4 implements Step 2 of the procedure, which propagates the STL formulae over a set of time intervals.

We now handle the approximated sets by extending the operations (1)–(3) on sets of time intervals. The procedures for the operations Invert, Join, Intersect, and $\mathsf{ShiftAll}_{\boldsymbol{t}}$ are described in Figure A·1 in the appendix. Note that, some operations cause ambiguities when handling non-canonical approximated time intervals. Such a situation is exemplified below. To avoid these ambiguities, our implementation results in an error once a resulting set becomes non-canonical.[†]

Example 7 *Consider the same timer system as in Example 5, i.e.* $\tilde{x}(t) := t$, *and the property* $\varphi := \mathsf{F}_{[0,\bar{t}]} \neg(x - 1 < 0 \vee 1 - x < 0)$, *where* $\bar{t} \in \mathbb{R}_{>0}$. *The subformula* $x - 1 < 0 \vee 1 - x < 0$ *is consistent at every time except at* $t = 1$, *therefore, the set of consistent time intervals is* $T := \{[0,1), [1, t_{\max})\}$. *Assume* T *is approximated with a non-canonical set* $\{([0], \mathsf{true}), ([0.95, 1.1], \mathsf{false}), ([0.9, 1.05], \mathsf{true})\}$.[††] *To verify* φ, *the procedures* Invert *and* $\mathsf{ShiftAll}_{[0,\bar{t}]}$ *should*

[†] To output unknown only due to the insufficient precision of numerical computation, the procedure should branch the process and proceed evaluation for both cases; implementation of such a procedure remains as a future work.

[††] The bound enclosures are usually very accurate, but kept large on this example to emphasize their impact.

be applied. However, as illustrated in Figure 4, we cannot decide whether the overlapping boundary intervals should be removed and expanded, or separated, and Propagate *results in an error.*[†††] *This ambiguous situation is avoided by using only canonical approximations. Note that, in some cases, this local ambiguity does not impact the global consistency. In the case of this example, if* $\overline{t} < 0.9$*, both scenarios lead to* false*, and forking the resolution process would be able to resolve the local ambiguity.*

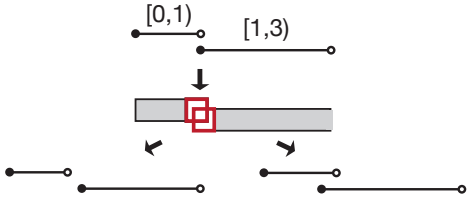

Fig. 4 Ambiguity caused by overlapping bounds

The following claims state that the procedures are closed in the canonical approximated sets, and the Propagate procedure is sound.

Lemma 1 *Let* $\boldsymbol{T}_1$ *and* $\boldsymbol{T}_2$ *be canonical approximated sets. If* $\boldsymbol{T}$ *results from* $\mathsf{Invert}(\boldsymbol{T}_1)$, $\mathsf{Join}(\boldsymbol{T}_1, \boldsymbol{T}_2)$, $\mathsf{Intersect}(\boldsymbol{T}_1, \boldsymbol{T}_2)$, *or* $\mathsf{ShiftAll}_t(\boldsymbol{T}_1, \boldsymbol{T}_2)$, *and if no* **error** *occurs in these procedures, then* $\boldsymbol{T}$ *is canonical.*

Proof. See Appendix A.1.

Theorem 2 (Soundness) *Consider an STL formula* φ *and a set* $\mathcal{T} = \{\boldsymbol{T}_p\}_{p \in AP_\varphi}$ *of approximated sets of time intervals that are consistent with atomic propositions. If* $\boldsymbol{T}_\varphi = \mathsf{Propagate}(\varphi, \mathcal{T})$*, then* $\boldsymbol{T}_\varphi$ *is an approximation of* T_φ*.*

Proof. See Appendix A.2.

Finally, we obtain $\boldsymbol{T}_\varphi$ and conclude that φ is valid if $\boldsymbol{s}_1$ is the smallest interval in $\boldsymbol{T}_\varphi$ and $\overline{s}_1 \leq 0 < \underline{s}'_1$, unsat if $\boldsymbol{T}_\varphi = \emptyset$ or $0 < \underline{s}_1$, or unknown if $0 \in [\underline{s}_1, \overline{s}_1)$. The computation is performed by ConsistentAtInitTime (see Algorithm 5 in the appendix).

6.4 Computational Cost

The time complexity of Propagate is bounded by the product of the size (i.e., the number of operators) of the considered STL formula φ and the cost of the procedures Invert, Join, and $\mathsf{ShiftAll}_t$ in Figure A·1. The complexity of the procedures on approximated sets is polynomial in the number of intersections of the signal and the atomic proposition bounds (see Appendix A.3). The number of iterations in MonitorAP is bounded by the product of the size of AP_φ and the maximum number of bounds detected for an atomic proposition, i.e., $\max_{p \in AP_\varphi} \#\boldsymbol{T}_p$.

The number of bounds detected depends on the oscillations of the ODE solution and the predicate bound. Although it can be very high in theory, it is usually quite small. In the generic case of non tangent intersection between the ODE solution and the predicate bound, the SearchZero procedure has a quadratic convergence and therefore a very low computational cost. The main computational cost of the method is therefore the validated simulation of the ODE. This cost is difficult to foresee: It highly depends on the ODE and the solver. For example, validated solvers are currently quite inefficient in solving stiff ODEs, and require to iterate many steps leading to a high computational cost. The complexity of Propagate is bounded by the product of the size (i.e., the number of operators) of the considered STL formula φ and the cost of the procedures in Figure A·1.

7. Experiments

We have implemented the proposed method and experimented on two examples to confirm the effectiveness of the method. Experiments were run on a 3.4GHz Intel Xeon processor with 16GB of RAM.

7.1 Implementation

Algorithms 1–5 were implemented in OCaml and C/C++. ODEs were solved by procedures in the CAPD library. The configurable parameters $t_{\min}$, ϵ, and θ correspond to the smallest integration step size that CAPD can take, the threshold used in Figure 3, and the threshold used in Inflate, respectively. In the experiments, these parameters were set as $t_{\min} := 10^{-14}$, $\epsilon := 10^{-14}$, and $\theta := 0.01$.

7.2 Verification of the Rotation System

We verified the system in Example 1 on four STL formulae. The specifications and results of this experiment are summarized in Table 1. The first column lists the STL formulae in which the bound τ of each G operator is parameterized and set to either $\tau := 100$ or $\tau := 10$. The column "$\#AP_\varphi$" represents the number of atomic propositions in each φ. In each verification, the parameter value u_1 was first randomly selected from $[-0.1, 0.1]$ and then modified to $u_1 := u_1 + \boldsymbol{u}_1$, where $\boldsymbol{u}_1$ was any of $[0]$, $[-10^{-6}, 10^{-6}]$, or $[-10^{-3}, 10^{-3}]$. The column "wid $\boldsymbol{u}_1$" indicates the interval used in each verification.

[†††] The property $\mathsf{F}_{[2,3]}\neg(x-1)^2 < 0$ is verified in the same way. The set T is consistent with the atomic proposition $(x-1)^2 < 0$. The verification will result in an error when SearchZero computes an enclosure of T at time 1.

Table 1 Experimental results (rotation)

φ	$\#AP_\varphi$	τ	wid $\boldsymbol{s}_1$	#valid	#unsat	#unknown	time
$\mathsf{G}_{[0,\tau]}\mathsf{F}_{[0,6.284]}\neg(x_2-1<0)$	1	100	0	507	493	0+0	0.51s
			$2\cdot10^{-6}$	483	508	9+0	0.52s
			$2\cdot10^{-3}$	0	462	538+0	–
		10	0	490	510	0+0	0.03s
			$2\cdot10^{-3}$	270	470	260+0	0.02s
$\mathsf{G}_{[0,\tau]}\mathsf{F}_{[0,6.284]}$ $(\neg(x_2-1<0)\wedge\mathsf{F}_{[0,3.142]}\neg(-x_2-1<0))$	2	100	0	485	515	0+0	1.1s
			$2\cdot10^{-6}$	353	505	26+116	1.03s
			$2\cdot10^{-3}$	0	463	537+0	–
		10	0	514	486	0+0	0.05s
			$2\cdot10^{-3}$	86	493	421+0	0.06s
$\mathsf{G}_{[0,\tau]}\mathsf{F}_{[0,6.284]}(\neg(x_2-1<0)\wedge$ $\mathsf{F}_{[0,1.571]}(\neg(-x_2<0)\wedge$ $\mathsf{F}_{[0,1.571]}(\neg(-x_2-1<0)\wedge$ $\mathsf{F}_{[0,1.571]}(-x_2<0)\ )))$	3	100	0	482	518	0+0	1.7s
			$2\cdot10^{-6}$	346	498	18+138	1.6s
			$2\cdot10^{-3}$	0	0	1000+0	–
		10	0	516	484	0+0	0.08s
			$2\cdot10^{-3}$	84	0	916+0	0.09s
$\mathsf{G}_{[0,\tau]}\mathsf{F}_{[0,6.284]}(\neg(x_2-1<0)\wedge$ $\mathsf{F}_{[0,0.786]}((x_2-0.707<0)\wedge\mathsf{F}_{[0,0.786]}(\neg(-x_2<0)\wedge$ $\mathsf{F}_{[0,0.786]}(\neg(-x_2-0.707<0)\wedge\mathsf{F}_{[0,0.786]}(\neg(-x_2-1<0)\wedge$ $\mathsf{F}_{[0,0.786]}((-x_2-0.707<0)\wedge\mathsf{F}_{[0,0.786]}((-x_2<0)\wedge$ $\mathsf{F}_{[0,0.786]}\neg(x_2-0.707<0)\ )))))))$	5	100	0	490	510	0+0	2.7s
			$2\cdot10^{-6}$	352	477	74+97	2.7s
			$2\cdot10^{-3}$	0	0	1000+0	–
		10	0	499	501	0+0	0.14
			$2\cdot10^{-3}$	0	0	1000+0	–

The considered STL properties are assumed to hold if $u_1 > 0$ and not to hold if $u_1 < 0$. Each STL property was verified for 1000 times. The columns "#valid", "#unsat", and "#unknown" list the numbers of runs resulting in each output; the "#unknown" outputs are separated with '+' according to whether it was caused by an error in the SearchZero algorithm or an error in the Propagate and ConsistentAtInitTime algorithms. The column "time" lists the average CPU time taken for a valid verification.

From the results, we can observe that the rates of inconclusive runs were related to the simulation lengths, the uncertainties in the parameter values, and the size of the formula φ. unknown results were generated by the interval Newton process in SearchZero and the undecidable situations in Propagate and ConsistentAtInitTime. In this experiment, verification failures increased as the value of u_1 approached 0 and the signal and boundary condition became close to tangent. When the parameter values were exact and wid $\boldsymbol{u}_1 = 0$, all the verifications succeeded even under near-singular conditions because the considered signals were always enclosed with tight intervals. As coarser intervals were appended to the parameter values and the simulation lengths became longer, the number of unknown results increased; meanwhile, the number of valid results decreased more rapidly than the number of unknown results because a valid verification required detecting a number of bounds for each atomic proposition. Any detection failure resulted in unknown.

The bottleneck of the verification process is the SearchZero algorithm that integrates ODEs and searches for boundary intervals. The number of calls to SearchZero depends on the size of AP_φ and the number of bounds as described in Section 6.4. Therefore, the runtime increased linearly in either the number of atomic propositions or the simulation length that should be proportional to the number of bounds. The cost of evaluation of the STL formulae seemed relatively small and not affecting the overall timings.

7.3 Verification of the Lorenz System

We verified the system in Example 2 on the following STL formula:

$$\mathsf{G}_{[0,15]}(\neg(-x_1-15<0)\Rightarrow \mathsf{F}_{[0.5,5]}\mathsf{G}_{[0,1]}((x_1-10)^2+(x_2-10)^2-150<0)) \quad (7)$$

In each verification, the parameters were set to exact values randomly selected from the domain. The signal (x_1, x_2) oscillates on either the positive or the negative side. According to the formula, when x_1 descends below -15, (x_1, x_2) moves into the disk $(x_1-10)^2+(x_2-10)^2<150$ after some duration in the interval $[0.5, 5]$ and remains in the disk for at least 1 time unit.

The experimental results are summarized in Table 2. As in Table 1, the columns (from left to right) represent the number of atomic propositions, the numbers of valid, unsat, and unknown verification results in 1000 runs, and the average CPU time for a valid verification.

Table 2 Experimental results (Lorenz)

$\#AP_\varphi$	#valid	#unsat	#unknown	time
2	566	413	21	9.2s

This experiment demonstrated that the proposed method can handle a chaotic system with a nonlinear atomic proposition. In such systems, non-validated numerical methods frequently output wrong results because of rounding errors, as shown in the next section.

As explained in Example 2, CAPD integration generated a coarse enclosure of the signal (around 25.8 time units), which introduced errors in the integration process X or the interval Newton process. These errors would account for the 21 unknown results in Table 2.

7.4 Comparison with Breach Toolbox

For comparative purposes, we ran the above problems on the Breach Toolbox [15] (built from commit ed1178c in the Mercurial repository), a tool for STL verification based on numerical computation with rounding errors. Breach can check the satisfiability and the *robustness*, which is quantified by a positive or negative real value based on the distance between a considered signal and the bound in the state space where the satisfaction of the STL property switches.

For the rotation system, when the parameter value u_1 approached 0 (specifically, at $u_1 := 0.001$), Breach returned unsat, whereas our implementation returned valid. This incorrect verification was implied by the low robustness value. In this example, the robustness was low for all parameter values because the initial part of the signal was close to the bounds of the atomic propositions.

For the Lorenz system, the numerical integration process of Breach yielded incorrect signals, as explained in Example 3; therefore, the verification results were unreliable. For example, when $u = (10, 28, 2.5)$, Breach reported an unsat verification of property (7), whereas our method returned certified valid.

Breach ran more quickly than our implementation: it required less than 0.01s for both problems.

8. Conclusions

We have presented a sound STL validation method for checking that all initialized signals satisfy the properties of a system. The proposed method detects a witness signal and verifies its unique existence using an interval-based ODE integration and an interval Newton method. The experimental results demonstrate the potential for the method as a practical tool.

In future work, we will improve our method and implementation to handle hybrid systems and large and uncertain initial values. Examples in a realistic setting should be demonstrated with the implementation.

Acknowledgments

This work was partially funded by JSPS (KAKENHI 25880008 and 15K15968).

References

[1] E. Plaku, L.E. Kavraki, and M.Y. Vardi, "Falsification of LTL Safety Properties in Hybrid Systems," TACAS, *LNCS* 5505, pp.368–382, 2009.
[2] T. Nghiem, S. Sankaranarayanan, G. Fainekos, F. Ivancic, A. Gupta, and G.J. Pappas, "Monte-Carlo Techniques for Falsification of Temporal Properties of Non-Linear Hybrid Systems," HSCC, pp.211–220, 2010.
[3] A. David, D. Du, K.G. Larsen, A. Legay, M. Mikučionis, D.B. Poulsen, and S. Sedwards, "Statistical Model Checking for Stochastic Hybrid Systems," Electronic Proceedings in Theoretical Computer Science, vol.92, pp.122–136, aug 2012.
[4] P. Zuliani, A. Platzer, and E.M. Clarke, "Bayesian statistical model checking with application to Stateflow/Simulink verification," Formal Methods in System Design, vol.43, no.2, pp.338–367, 2013.
[5] A. Eggers, M. Franzle, and C. Herde, "SAT Modulo ODE : A Direct SAT Approach to Hybrid Systems," ATVA, *LNCS* 5311, no.1, pp.171–185, 2008.
[6] P. Collins and A. Goldsztejn, "The Reach-and-Evolve Algorithm for Reachability Analysis of Nonlinear Dynamical Systems," Electronic Notes in Theoretical Computer Science, vol.223, no.639, pp.87–102, dec 2008.
[7] N. Ramdani and N.S. Nedialkov, "Computing reachable sets for uncertain nonlinear hybrid systems using interval constraint-propagation techniques," Nonlinear Analysis: Hybrid Systems, vol.5, no.2, pp.149–162, may 2011.
[8] D. Ishii, K. Ueda, and H. Hosobe, "An interval-based SAT modulo ODE solver for model checking nonlinear hybrid systems," International Journal on Software Tools for Technology Transfer (STTT), vol.13, no.5, pp.449–461, 2011.
[9] X. Chen, E. Abraham, and S. Sankaranarayanan, "Taylor Model Flowpipe Construction for Non-linear Hybrid Systems," IEEE Real-Time Systems Symposium, pp.183–192, 2012.
[10] S. Gao and E.M. Clarke, "Satisfiability Modulo ODEs," FMCAD, pp.105–112, 2013.
[11] S. Gao, J. Avigad, and E.M. Clarke, "Delta-Decidability over the Reals," LICS, pp.305–314, 2012.
[12] O. Maler and D. Nickovic, "Monitoring Temporal Properties of Continuous Signals," FORMATS, *LNCS* 3253, pp.152–166, 2004.
[13] G. Fainekos, A. Girard, and G. Pappas, "Temporal logic verification using simulation," FORMATS, *LNCS* 4202, vol.4202, pp.171–186, 2006.
[14] A. Donzé and O. Maler, "Robust Satisfaction of Temporal Logic over Real-Valued Signals," FORMATS, *LNCS* 6246, pp.92–106, 2010.
[15] A. Donzé, "Breach, a toolbox for verification and parameter synthesis of hybrid systems," CAV, *LNCS* 6174, pp.167–170, 2010.
[16] E. Goubault, O. Mullier, and M. Kieffer, "Inner Approximated Reachability Analysis," HSCC, pp.163–172, 2014.
[17] R. Alur, T. Feder, and T.A. Henzinger, "The Benefits of Relaxing Punctuality," Journal of the ACM, vol.43, no.1, pp.116–146, 1996.
[18] B. Shultz and B.J. Kuipers, "Proving properties of continuous systems : qualitative simulation and temporal logic," Artificial Intelligence, vol.92, no.96, pp.91–129, 1997.
[19] Q. Wang, P. Zuliani, S. Kong, S. Gao, and E. Clarke, "SReach : Combining Statistical Tests and Bounded Model Checking for Nonlinear Hybrid Systems with Parametric Uncertainty," tech. rep., Carnegie Mellon University, 2014.
[20] A. Podelski and S. Wagner, "Model Checking of Hybrid Systems : From Reachability towards Stability," HSCC, *LNCS* 3927, pp.507–521, 2006.
[21] A. Cimatti, A. Griggio, S. Mover, and S. Tonetta, "Verifying LTL Properties of Hybrid Systems with K-Liveness," CAV, *LNCS* 8559, pp.424–440, 2014.

[22] R.E. Moore, Interval Analysis, Prentice-Hall, 1966.
[23] A. Neumaier, Interval Methods for Systems of Equations, Cambridge University Press, 1990.
[24] N.S. Nedialkov, "VNODE-LP — A Validated Solver for Initial Value Problems in Ordinary Differential Equations," tech. rep., McMaster University, 2006.
[25] A. Goldsztejn and L. Jaulin, "Inner approximation of the range of vector-valued functions," Reliable Computing, vol.14, pp.1–23, 2010.

Appendix: Omitted Procedures and Proofs

Procedures of the operations on approximated sets are specified in Figure A·1. Invert, Join, Intersect, and $\mathsf{ShiftAll}_{\boldsymbol{t}}$ implement the operations in Step 2 of Section 5.2 as procedures that modify the set of boundary intervals. $\mathsf{ShiftAll}_{\boldsymbol{t}}$ consists of sub-procedures ShiftPairs and ShiftElem; ShiftPairs computes the intersections and back-shifting pairwise ($\mathsf{Pairs}(\boldsymbol{T})$ enumerates approximations of time intervals in $\boldsymbol{T}$); ShiftElem applies the back-shifting. The results of the procedures may become non-canonical, so Normalize is applied at last to make them canonical.

ConsistentAtInitTime is implemented in Algorithm 5. An input $\boldsymbol{T}_\varphi$ is either $\boldsymbol{T}_{\mathsf{true}}$, $\emptyset$, or an approximated set; in the last case, the algorithm picks an earliest approximation with GetFirstElem, and checks whether it contains 0 or not.

Algorithm 5 ConsistentAtInitTime algorithm

Input: $\boldsymbol{T}_\varphi$
Output: valid, unsat, or unknown
1: **if** $\boldsymbol{T}_\varphi = \boldsymbol{T}_{\mathsf{true}}$ **then**
2: **return** valid
3: **else if** $\boldsymbol{T}_\varphi = \emptyset$ **then**
4: **return** unsat
5: **else**
6: $(\boldsymbol{s}, \mathsf{true}) := \mathsf{GetFirstElem}(\boldsymbol{T}_\varphi)$
7: **if** $\overline{s} \leq 0$ **then**
8: **return** valid
9: **else if** $\underline{s} > 0$ **then**
10: **return** unsat
11: **else** $\{0 \in \boldsymbol{s}\}$
12: **return** unknown
13: **end if**
14: **end if**

A.1 Proof of Lemma 4

We check that each condition of a canonical approximation (Definition 6) is assured by Normalize, the last sub-process in each procedure:

- During the propagation process, polarity alternation might be inhibited by Join, Intersect, or $\mathsf{ShiftAll}_{\boldsymbol{t}}$, which locates a boundary interval inside another consistent time interval. These *embedded* bounds are removed by N_1. An embodiment can be determined by checking the difference between the numbers of lower and upper bounds in the past since the smallest elements in $\boldsymbol{T}_1$ and $\boldsymbol{T}_2$ are always the lower-bound enclosures.
- The upper bound of each time interval $\boldsymbol{s}$ in $\boldsymbol{T}$ becomes non-negative because elements with non-positive upper bounds are filtered out by N_2.
- No two elements of $\boldsymbol{T}$ overlap because an overlapping pair with opposite polarity results in an error (the second branch of Normalize) and an overlap with the same polarity is joined (N_3); thus, the elements in $\boldsymbol{T}$ can be sorted.
- N_4 assures that the polarity value of the smallest element is true. □

A.2 Proof of Theorem 2

We perform a structural induction based on the STL formulae.

For the base case $\varphi = p \in AP_\varphi$, $\boldsymbol{T}_p$ exists in $\mathcal{T}$.

For the inductive step, consider STL formulae φ_1 and φ_2, and assume as the inductive hypothesis that we have canonical approximated sets $\boldsymbol{T}_{\varphi_1}$ and $\boldsymbol{T}_{\varphi_2}$ of T_{φ_1} and T_{φ_2}, respectively. We show that Propagate computes the approximated set properly for each formula constructed from φ_1 and φ_2.

When $\varphi = \neg\varphi_1$, the polarity of each bound of $\boldsymbol{T}_{\varphi_1}$ is switched by Invert to obtain an approximated set for the complementary time intervals, which is sound regarding the operation (1) in Step 2 of Section 5.2. Then, Normalize is applied to canonicalize the result; it will append or remove the smallest bound. Let $(\boldsymbol{s}, \mathsf{false})$ be the smallest element in a result of polarity inversion. We confirm that Normalize is sound in a case analysis:

- if $\boldsymbol{s}$ is non-empty and $\boldsymbol{s} \ni 0$, the computation results in an error (the first branch of Normalize);
- if $\underline{s} > 0$, the element remains and the element $([0], \mathsf{true})$ is appended by N_3;
- if $\boldsymbol{s} = [0]$, the element is removed by N_2.

When $\varphi = \varphi_1 \vee \varphi_2$, $\boldsymbol{T}_{\varphi_1}$ and $\boldsymbol{T}_{\varphi_2}$ are modified by Join, which joins the elements of both approximated sets; a result might be a non-canonical set when two approximated time intervals from $\boldsymbol{T}_{\varphi_1}$ and $\boldsymbol{T}_{\varphi_2}$ overlap. Then, Normalize is applied to unify two overlapping approximations so that the result becomes a sound approximated set with respect to the operation (2). When two approximated time intervals $((\boldsymbol{s}_1, \mathsf{true}), (\boldsymbol{s}'_1, \mathsf{false}))$ and $((\boldsymbol{s}_2, \mathsf{true}), (\boldsymbol{s}'_2, \mathsf{false}))$ overlap, the boundary interval (e.g., $\boldsymbol{s}_1$) either (i) overlaps with another boundary interval, (ii) is included in the inner approximation $(\overline{s}_2, \underline{s}'_2)$, or (iii) is excluded from the outer approximation $[\underline{s}_2, \overline{s}'_2]$. We confirm the soundness of Normalize

$$\mathsf{Invert}(\boldsymbol{T}) := \begin{cases} \emptyset & \text{if } \boldsymbol{T} = \boldsymbol{T}_{\mathsf{true}} \\ \boldsymbol{T}_{\mathsf{true}} & \text{if } \boldsymbol{T} = \emptyset \\ \mathsf{Normalize}(\ \{(\boldsymbol{s}, \neg b) \mid (\boldsymbol{s}, b) \in \boldsymbol{T}\}\ ) & \text{otherwise} \end{cases}$$

$$\mathsf{Join}(\boldsymbol{T}_1, \boldsymbol{T}_2) := \begin{cases} \boldsymbol{T}_{\mathsf{true}} & \text{if } \boldsymbol{T}_1 = \boldsymbol{T}_{\mathsf{true}} \vee \boldsymbol{T}_2 = \boldsymbol{T}_{\mathsf{true}} \\ \mathsf{Normalize}(\ \boldsymbol{T}_1 \cup \boldsymbol{T}_2\ ) & \text{otherwise} \end{cases}$$

$$\mathsf{Intersect}(\boldsymbol{T}_1, \boldsymbol{T}_2) := \mathsf{Invert}(\ \mathsf{Join}(\mathsf{Invert}(\boldsymbol{T}_1), \mathsf{Invert}(\boldsymbol{T}_2))\ )$$

$$\mathsf{ShiftAll}_{\boldsymbol{t}}(\boldsymbol{T}_1, \boldsymbol{T}_2) := \begin{cases} \emptyset & \text{if } \boldsymbol{T}_1 = \emptyset \vee \boldsymbol{T}_2 = \emptyset \\ \mathsf{Normalize}(\ \mathsf{ShiftPairs}_{\boldsymbol{t}}(\boldsymbol{T}_1, \boldsymbol{T}_2)\ ) & \text{otherwise} \end{cases}$$

$$\mathsf{ShiftPairs}_{\boldsymbol{t}}(\boldsymbol{T}_1, \boldsymbol{T}_2) := \begin{cases} \{\mathsf{ShiftElem}_{\boldsymbol{t}}(\boldsymbol{P}_2) \mid \boldsymbol{P}_2 \in \mathsf{Pairs}(\boldsymbol{T}_2)\} & \text{if } \boldsymbol{T}_1 = \boldsymbol{T}_{\mathsf{true}} \\ \{\mathsf{Intersect}(\ \mathsf{ShiftElem}_{\boldsymbol{t}}(\boldsymbol{P}_1),\ \boldsymbol{P}_1\ ) \mid \boldsymbol{P}_1 \in \mathsf{Pairs}(\boldsymbol{T}_1)\} & \text{if } \boldsymbol{T}_2 = \boldsymbol{T}_{\mathsf{true}} \\ \{\mathsf{Intersect}(\ \mathsf{ShiftElem}_{\boldsymbol{t}}(\mathsf{Intersect}(\boldsymbol{P}_1, \boldsymbol{P}_2)),\ \boldsymbol{P}_1\ ) \mid \boldsymbol{P}_1 \in \mathsf{Pairs}(\boldsymbol{T}_1), \boldsymbol{P}_2 \in \mathsf{Pairs}(\boldsymbol{T}_2)\} & \text{otherwise} \end{cases}$$

$$\mathsf{ShiftElem}_{\boldsymbol{t}}(\boldsymbol{T}) := \mathsf{Normalize}(\ \{(\boldsymbol{s} - \underline{t}, \mathsf{true}) \mid \exists (\boldsymbol{s}, \mathsf{true}) \in \boldsymbol{T}\} \cup \{(\boldsymbol{s} - \overline{t}, \mathsf{false}) \mid (\boldsymbol{s}, \mathsf{false}) \in \boldsymbol{T}\}\ )$$

$$\mathsf{Normalize}(\boldsymbol{T}) := \begin{cases} \textbf{error} & \text{if } \exists (\boldsymbol{s}, \mathsf{false}) \in \boldsymbol{T}\ \boldsymbol{s} \neq [0] \wedge \boldsymbol{s} \ni 0 \\ \textbf{error} & \text{if } \exists (\boldsymbol{s}, b), (\boldsymbol{s}', \neg b) \in \boldsymbol{T}\ \boldsymbol{s} \cap \boldsymbol{s}' \neq \emptyset \\ \mathsf{N}_4(\mathsf{N}_3(\mathsf{N}_2(\mathsf{N}_1(\boldsymbol{T})))) & \text{otherwise} \end{cases}$$

$$\mathsf{N}_1(\boldsymbol{T}) := \{(\boldsymbol{s}, \mathsf{true}) \in \boldsymbol{T} \mid \#\{(\boldsymbol{s}', \mathsf{true}) \in \boldsymbol{T} \mid \overline{s}' < \underline{s}\} - \#\{(\boldsymbol{s}'', \mathsf{false}) \in \boldsymbol{T} \mid \overline{s}'' < \underline{s}\} < 1\} \cup \{(\boldsymbol{s}, \mathsf{false}) \in \boldsymbol{T} \mid \#\{(\boldsymbol{s}', \mathsf{true}) \in \boldsymbol{T} \mid \overline{s}' < \underline{s}\} - \#\{(\boldsymbol{s}'', \mathsf{false}) \in \boldsymbol{T} \mid \overline{s}'' < \underline{s}\} < 2\}$$

$$\mathsf{N}_2(\boldsymbol{T}) := \begin{cases} \boldsymbol{T}_{\mathsf{true}} & \text{if } \max \boldsymbol{T} = (\boldsymbol{s}, \mathsf{true}) \text{ such that } \overline{s} \leq 0 \\ \{(\boldsymbol{s}, b) \in \boldsymbol{T} \mid \overline{s} > 0\} & \text{otherwise} \end{cases}$$

$$\mathsf{N}_3(\boldsymbol{T}) := \{(\boldsymbol{s}, b) \in \boldsymbol{T} \mid \forall (\boldsymbol{s}', b) \in \boldsymbol{T}\ \boldsymbol{s} \cap \boldsymbol{s}' = \emptyset\} \cup \{(\boldsymbol{s} \cup \boldsymbol{s}', b) \mid \exists (\boldsymbol{s}, b), (\boldsymbol{s}', b) \in \boldsymbol{T}\ \boldsymbol{s} \cap \boldsymbol{s}' \neq \emptyset\}$$

$$\mathsf{N}_4(\boldsymbol{T}) := \begin{cases} \boldsymbol{T} \cup \{([0], \mathsf{true})\} & \text{if } (\boldsymbol{t}, \mathsf{false}) = \min \boldsymbol{T} \\ \boldsymbol{T} & \text{otherwise} \end{cases}$$

Fig. A· 1 Procedures for approximated sets of consistent time intervals

in another case analysis:

- in case (i), the bound is removed by the second branch of Normalize and by N_3;
- in case (ii), the bound is removed by N_1;
- in case (iii), the bound remains since it should be the bound of the joined time interval.

When $\varphi = \varphi_1 \mathsf{U}_{\boldsymbol{t}} \varphi_2$, $\boldsymbol{T}_{\varphi_1}$ and $\boldsymbol{T}_{\varphi_2}$ are modified by $\mathsf{ShiftAll}_{\boldsymbol{t}}$, which applies Intersect, $\mathsf{ShiftPairs}_{\boldsymbol{t}}$ and $\mathsf{ShiftElem}_{\boldsymbol{t}}$, those implement the operation (3). The soundness of Intersect with respect to the set intersection is evident because this procedure simply implements the set operation $(\boldsymbol{T}_1 \setminus \mathbb{R}_{\geq 0} \cup \boldsymbol{T}_2 \setminus \mathbb{R}_{\geq 0}) \setminus \mathbb{R}_{\geq 0}$. $\mathsf{ShiftPairs}_{\boldsymbol{t}}$ exhaustively applies $\mathsf{ShiftElem}_{\boldsymbol{t}}$ to each pair of boundary enclosures in $\boldsymbol{T}_1$ and $\boldsymbol{T}_2$. $\mathsf{ShiftElem}_{\boldsymbol{t}}$ translates the lower and upper bounds, according to the operation (3); this procedure is sound because an interval enclosure is assumed for each bound of the consistent time intervals. Normalize, then, resolves the overlaps and closes the lowest bound as in the case of $\varphi_1 \vee \varphi_2$. □

A.3 Computational Complexity of the Operations on Approximated Sets

Let $\#\boldsymbol{T}$ be the number of elements in $\boldsymbol{T}$; if the bounds appear uniformly in a simulation, $\#\boldsymbol{T}$ is proportional to $\|\varphi\|$; in other words, $\#\boldsymbol{T}$ is bounded by $\|\varphi\|/\epsilon^*$ where ϵ^* is the precision of the floating-point numbers. The complexity of Normalize is bounded by $O(\#\boldsymbol{T}^2)$ since the complexities of N_1, N_2, N_3, and N_4 are $O(\#\boldsymbol{T}^2)$, $O(\#\boldsymbol{T})$, $O(\#\boldsymbol{T})$, and $O(1)$, respectively. Without the Normalize process, the complexities of Invert and Join are $O(\#\boldsymbol{T})$ and $O(1)$, respectively; together with Normalize, their complexities are $O(\#\boldsymbol{T}^2)$. The complexity of $\mathsf{ShiftAll}_{\boldsymbol{t}}$ is $O(\#\boldsymbol{T}^4)$ (let $\#\boldsymbol{T}$ be the larger cardinality for $\boldsymbol{T}_1$ or $\boldsymbol{T}_2$) since the complexities of ShiftElem and ShiftPairs are $O(\#\boldsymbol{T}^2)$ and $O(\#\boldsymbol{T}^2 \cdot \#\boldsymbol{T}^2)$, respectively.

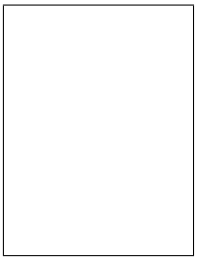

Daisuke Ishii recieved B.Eng, M.Eng, and Ph.D. degrees in computer science from Waseda University, Tokyo, Japan in 2001, 2003, and 2010, respectively. He was a research fellow at INRIA/LINA in France, from 2010 to 2011, and a research fellow of the JSPS at National Institute of Informatics from 2011 to 2013. He is currently an Assistant Professor at Tokyo Institute of Technology. His research interests include interval analysis and formal methods for hybrid systems.
Tokyo Institute of Technology, Department of Computer Science, 2-12-1-W8-67, Ookayama, Meguro-ku, Tokyo, 152-8550 Japan.

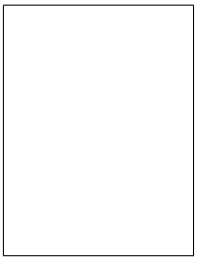

Naoki Yonezaki Naoki Yonezaki currently is Visiting Professor of Open University of Japan and Emeritus Professor of Tokyo Institute of Technology. He has been Professor of Tokyo Institute of Technology since 1991. He was also Professor of Japan Advanced Institute of Science and Technology from1991 till 1995. His research interests include verification of software specification, verification of security and formal approach to system biology. He is a member of IEICE, IPSJ, JSSST, JSAI, ACM and EATCS. He awarded to a fellowship from JSSST in 2008.
Open University of Japan, 2-11 Wakaba, Mihama-ku, Chiba 261-8586 Japan.

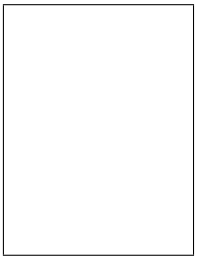

Alexandre Goldsztejn After receiving his Ph.D. from the University of Nice in 2005, he has spent one year as a postdoctoral fellow in USA, in the University of Central Arkansas and in the University of California Irvine. He was then granted a tenured research associate position at CNRS, where he continued his researches on numerical constraint programming, with emphasis on quantified constraints and positive dimensional manifolds, global optimization and dynamical systems.
IRCCyN – Ecole Centrale de Nantes, 1, rue de la Noë, BP 92101, 44321 Nantes Cedex 3, France.